

Stochastic Mapping Using Forward Look Sonar

John J. Leonard¹, Robert N. Carpenter², and Hans Jacob S. Feder¹

¹Massachusetts Institute of Technology, 77 Mass Ave., Cambridge, MA 02139

²Naval Undersea Warfare Center, Newport, RI 02841

jleonard@mit.edu, carpenterrn@code80.npt.nuwc.navy.mil, hjfeder@alum.mit.edu

Abstract

This paper investigates the problem of concurrent mapping and localization (CML) using forward look sonar data. Results are presented from processing of an oceanic data set from an 87 kHz US Navy forward look imaging sonar using the stochastic mapping method for CML. The goal is to detect objects on the seabed, map their locations, and concurrently compute an improved trajectory for the vehicle. The resulting trajectory is compared with position estimates computed with an inertial navigation system and Doppler velocity sonar. The results demonstrate the potential of concurrent mapping and localization algorithms to satisfy the navigation requirements of undersea vehicles equipped with forward look sonar.

1 Introduction

Navigation is a critical issue in the development of new autonomous underwater vehicle (AUV) technology [8]. Good positioning information is vital for the safe execution of an AUV mission and for effective interpretation of the data acquired by the AUV [14, 16]. Current methods for AUV navigation suffer several shortcomings. Dead reckoning and inertial navigation systems (INS) are subject to external disturbances and uncorrectable drift. Measurements from Doppler velocity sonars can be used to achieve higher precision, but position error still grows without bound. To achieve bounded errors, current AUV systems rely on networks of acoustic transponders or surfacing for GPS resets. The goal of concurrent mapping and localization (CML) is to enable an AUV to build a map of an unknown environment and concurrently use that map for positioning. CML has the potential to enable long-term missions with bounded navigation errors without reliance on acoustic beacons, a priori maps, or surfacing for GPS resets [13, 12].

Report Documentation Page			Form Approved OMB No. 0704-0188		
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 2001	2. REPORT TYPE		3. DATES COVERED 00-00-2001 to 00-00-2001		
4. TITLE AND SUBTITLE Stochastic Mapping Using Forward Look Sonar			5a. CONTRACT NUMBER		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S)			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Massachusetts Institute of Technology,Cambridge,MA,02139			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSOR/MONITOR'S ACRONYM(S)		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S)		
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES In Robotica, pages 467-480. 2001. U.S. Government or Federal Rights License					
14. ABSTRACT This paper investigates the problem of concurrent mapping and localization (CML) using forward look sonar data. Results are presented from processing of an oceanic data set from an 87 kHz US Navy forward look imaging sonar using the stochastic mapping method for CML. The goal is to detect objects on the seabed, map their locations, and concurrently compute an improved trajectory for the vehicle. The resulting trajectory is compared with position estimates computed with an inertial navigation system and Doppler velocity sonar. The results demonstrate the potential of concurrent mapping and localization algorithms to satisfy the navigation requirements of undersea vehicles equipped with forward look sonar.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 25	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

For several decades, the integration of mapping and navigation has been an important goal in the robotics research community [3, 6, 29, 21, 27, 1]. The problem is difficult because of the combination of sensor noise, data association ambiguity, navigation error, and modeling uncertainty. Considerable progress has been made in recent years, with new insights into the structure of the problem [9, 5, 10] and new approaches that have provided compelling experimental demonstrations [30, 15].

Two of the most difficult issues in CML are choosing a representation of the environment and coping with uncertainty in sensor data. Most successful implementations of CML have been performed in relatively structured indoor environments, with sensor data of fairly high quality. For example, the SICK laser scanner provides accurate range data that supports fast matching of scans taken from different robot positions [30, 15]. Feature-based approaches to CML have primarily been tested in situations with relatively distinct features, such as radar reflectors [9] or simple geometric objects [12]. In more complex environments, the issues of reliably extracting features from sensor data and performing data association become much more difficult.

In this paper, we investigate some of the issues encountered in testing of CML algorithms for use on AUVs. The stochastic mapping method for CML [28, 21] is applied to data acquired in Naragansett Bay, RI with the US Navy high resolution array (HRA) forward look imaging sonar [22, 20, 4]. Features are extracted from the sonar images, and supplied as input to the CML algorithm. The goal is to detect objects on the seabed, map their locations, and concurrently compute an improved trajectory for the vehicle. Many of the bottom objects encountered in the data are extended features rather than distinct points, making the data association task very challenging.

The structure of this paper is as follows: Section 2 reviews the stochastic mapping algorithm. Results from processing of a testing tank sonar data set are used to illustrate the method’s performance in a simple environment. Section 3 provides details on the HRA data acquisition and Section 4 describes the sonar signal and image processing operations that were performed on the raw data. Section 5 applies stochastic mapping to two of the HRA data sets. Finally, Section 6 summarizes our results and draws conclusions.

2 Stochastic mapping

The feature-based CML algorithm used to process the data is an augmented form of stochastic mapping, a Kalman filter-based method for CML first proposed by Smith, Self, and Cheeseman [28] and first implemented by Moutarlier and Chatila [21]. The method assumes that point-like features are

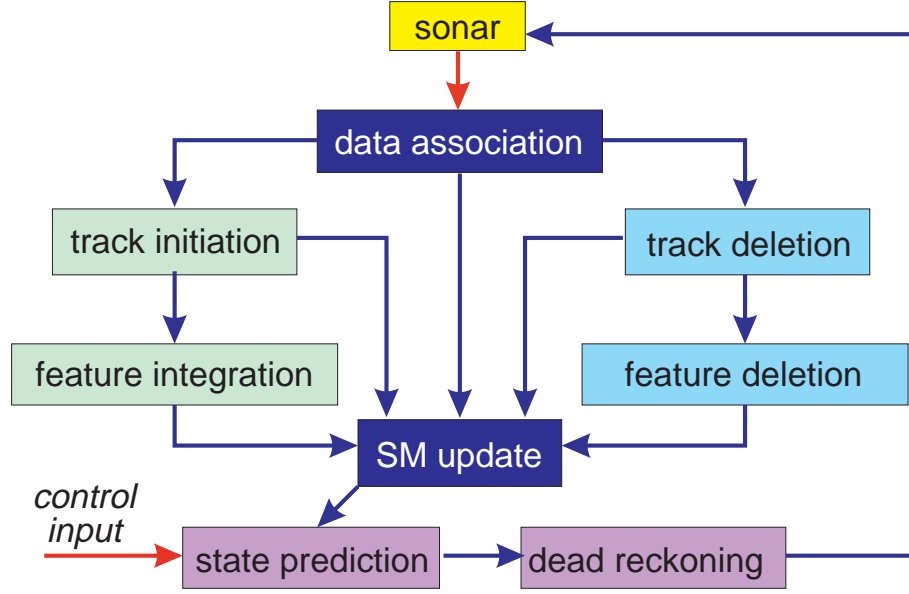


Figure 1: The stochastic mapping method for CML.

present in the environment and can be effectively detected and tracked.

Figure 1 provides an overview of the processing stages in stochastic mapping. In our implementation, the AUV senses features in the environment through range and bearing measurements. Estimates of feature locations from different vantage points are fused together to build a map, and concurrently, the feature location estimates are used to update the robot position. The robot and the map are represented by a single state vector, $\mathbf{x}(k)$, with an associated estimate error covariance $\mathbf{P}(k)$ at each time step. As new features are added, $\mathbf{x}(k)$ and $\mathbf{P}(k)$ increase in size. Measurements are taken every $t = kT$ seconds, where T is a constant period and k is a discrete time index.

We suppose that there are n features in the environment, and that they are static. The true state at discrete time k is designated by $\mathbf{x}(k) = [\mathbf{x}_r(k)^T \mathbf{x}_f(k)^T]^T$, where $\mathbf{x}_r(k)$ represents the location of the robot, and $\mathbf{x}_f(k)^T = [\mathbf{x}_1(k)^T \dots \mathbf{x}_n(k)^T]^T$ represent the locations of the environmental features. Let $\mathbf{z}(k)$ designate the sensor measurements obtained at time k , and \mathbf{Z}^k designate the set of all measurements obtained from time step 0 through time step k .

Stochastic mapping algorithms for CML use the extended Kalman filter to recursively compute a state estimate $\hat{\mathbf{x}}(k) = [\hat{\mathbf{x}}_r(k)^T \hat{\mathbf{x}}_f(k)^T]^T$ at each discrete time step k , where $\hat{\mathbf{x}}_r(k)^T$ and $\hat{\mathbf{x}}_f(k)^T = [\hat{\mathbf{x}}_1(k)^T \dots \hat{\mathbf{x}}_n(k)^T]^T$ are the estimated robot and feature locations, respectively. Based on assumptions about linearization and data association, this estimate is the approximate conditional mean of $p(\mathbf{x}(k)|\mathbf{Z}^k)$:

$$\hat{\mathbf{x}}(k) \approx E[p(\mathbf{x}(k)|\mathbf{Z}^k)].$$

Associated with this state vector is an estimated error covariance, $\mathbf{P}(k)$, which represents the errors in the robot and feature locations, and the cross-correlations between these states:

$$\mathbf{P}(k) = \begin{bmatrix} \mathbf{P}_{rr}(k) & \mathbf{P}_{rf}(k) \\ \mathbf{P}_{fr}(k) & \mathbf{P}_{ff}(k) \end{bmatrix} = \begin{bmatrix} \mathbf{P}_{rr}(k) & \mathbf{P}_{r1}(k) & \cdots & \mathbf{P}_{rn}(k) \\ \mathbf{P}_{1r}(k) & \mathbf{P}_{11}(k) & \cdots & \mathbf{P}_{1n}(k) \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{P}_{nr}(k) & \mathbf{P}_{n1}(k) & \cdots & \mathbf{P}_{nN}(k) \end{bmatrix}. \quad (1)$$

Caution must be used in any extended Kalman filter implementation, because $\mathbf{P}(k)$ is only an approximation to $\mathcal{P}(k)$, the actual mean squared error of the estimate at time k .

We denote the vehicle's state by $\mathbf{x}_r = [x_r \ y_r \ \phi \ v]^T$ to represent the vehicle's east position, north position, heading, and speed, respectively. The state of feature i is represented by $\mathbf{x}_i = [x_i \ y_i]^T$. The dynamic model used in the algorithm simulates an AUV equipped with control surfaces and a single aft thruster for propulsion, moving at a nominal forward speed of 2.5 m/s. The control input $\mathbf{u}(k)$ to the vehicle is given by a change in heading, $\delta\phi$, and speed, δv , of the vehicle to model changes in rudder angle and forward thrust, that is, $\mathbf{u}(k) = [\delta\phi \ \delta v]^T$. Thus, the dynamic model of the AUV, $\mathbf{f}()$, is given by

$$\mathbf{x}[k+1] = \mathbf{f}(\mathbf{x}(k), \mathbf{u}(k)) + \mathbf{d}_x(\mathbf{u}(k)), \quad (2)$$

where $\mathbf{d}_x(\mathbf{u}(k))$ is a white, Gaussian random process independent of $\mathbf{x}(0)$, with magnitude dependent on the control input $\mathbf{u}(k)$.

The observation model $\mathbf{h}()$ for the system is given by

$$\mathbf{z}(k) = \mathbf{h}(\mathbf{x}(k)) + \mathbf{d}_z, \quad (3)$$

where $\mathbf{z}(k)$ is the observation vector of range and bearing measurements. The observation model, $\mathbf{h}()$, defines the (nonlinear) coordinate transformation from state to observation coordinates. The stochastic process \mathbf{d}_z , is assumed to be white, Gaussian, and independent of $\hat{\mathbf{x}}(0)$ and \mathbf{d}_x , and has covariance \mathbf{R} . Given these assumptions, an extended Kalman filter (EKF) is employed to estimate the state $\hat{\mathbf{x}}$ and covariance \mathbf{P} given the measurements. Our implementation is similar to that of Smith *et al.*, and the reader is referred to [28] for a detailed derivation.

Data association refers to the task of determining the origins of measurements, and can be a critical issue in the implementation of stochastic mapping with real data [7]. Useful insights into the data association problem in CML are provided by the field of multiple target tracking [2]. For this reason, we refer to environmental feature as “targets”, and state estimates for environmental features as

“tracks” [17]. In multiple target tracking, a wide range of data association methods have been developed, with greatly varying computational requirements [2, 23]. In CML, the addition of vehicle navigation uncertainty greatly increases the complexity of the data association problem [25, 26].

In the results presented in this paper, the assignment of measurements to existing features is carried out by nearest neighbor gating [2]. For each feature in the vehicle state vector, predicted range and angle measurements are generated and are compared against the actual measurements using a weighted statistical distance in measurement space. That is, for all measurements \mathbf{z}_j that can potentially be associated with feature i , the innovation, $\boldsymbol{\nu}_{ij}$, and the innovation matrix, \mathbf{S}_{ij} , are constructed and the closest measurement within the gate defined by the Mahalanobis distance

$$\boldsymbol{\nu}_{ij}^T \mathbf{S}_{ij}^{-1} \boldsymbol{\nu}_{ij} \leq \gamma, \quad (4)$$

is considered the most likely measurement of that feature. Under the assumption that the measurement model noise is Gaussian, Equation (4) will have a χ^2 -distribution [2]. Thus, for a system of 2 degrees of freedom (the x and y positions of the feature), the true measurement of the feature, if detected, would fall within a $\gamma = 9$ gate with 99% probability. Equation (4) can be considered to define a “validation region” or “validation gate” for feature i . If measurement \mathbf{z}_j is the only measurement at a given time step that satisfies Equation (4), it is said to have “gated” with feature i . (The term “gate” originates from the radar term “range gate” [24].)

Measurements that do not gate with any existing feature become candidates for the initialization of new features in a process of delayed nearest neighbor track initiation. In this approach, all measurements that have not been associated with a track (feature) over the last N time steps are stored. At each time step, a search for clusters with more than $M \leq N$ measurements are performed. When such a cluster is found, a new track is initiated and the feature is integrated into the system state vector and system error covariance.

Figures 2 to 4 illustrate the performance of the stochastic mapping method in a simple environment using sonar data from a testing tank. (More detail is provided in Feder [11].) The experiment was designed to simulate an AUV performing CML with a sensor such as the HRA, with the AUV’s parameters scaled down by a factor of 100. The algorithm parameters for the tank experiment are shown in Table 1.

Figure 2 (top left) shows a photograph of the testing tank and robotic positioning system used to acquire the data. Figure 2 (top right) shows a typical raw sonar scan obtained with the sensor. Figure 2 (bottom left) shows the survey path of the vehicle and the position of all the objects for the experiment.

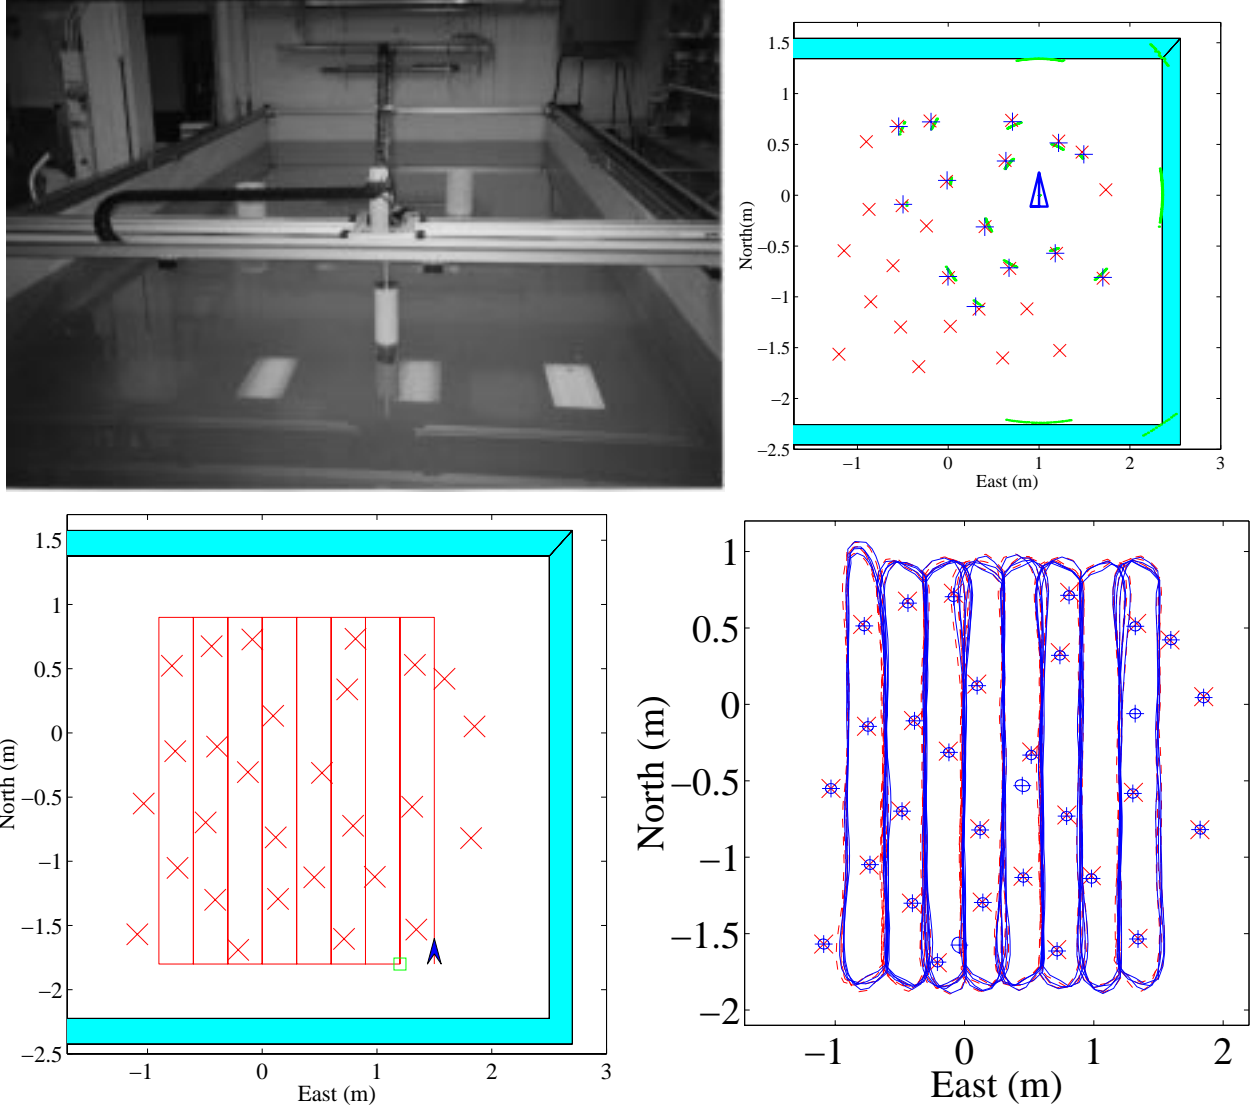


Figure 2: (Top left) Testing tank and robotic positioning system used to acquire the data. (Top right) A 360° scan of the tank with the Panametrics sonar. The location of the sonar is marked by a triangle. Individual sonar returns are marked by a dot. The true location of a feature is marked by 'x'. The tank walls are shaded. (Bottom left) The desired path of the sensor is shown by a solid line and the true feature positions are marked by crosses. The initial position and direction of the AUV is marked by an arrow. The end of the path before it is repeated is marked by a square. (Bottom right) The estimated path is shown by a solid line and the actual path is shown by a dashed line. The true feature location are marked by 'x' and the estimated feature locations by '+'. 3σ error ellipse are also shown for the feature estimates.

Table 1: Tank experiment parameters.

sampling period, T	1 sec.
maximum sonar range	200 cm
sonar coverage angle	$\pm 40^\circ$
range measurement standard deviation	2 cm
bearing measurement standard deviation	5°
feature probability of detection, P_D	0.90
vehicle cruise speed, V_c	10 cm/T
speed process standard deviation	5% of δv
heading process standard deviation	2.0°
dead reckoning speed standard deviation	4.5 cm/T
dead reckoning heading standard deviation	3.0°
initial position uncertainty std. dev.	1.0 cm
initial heading uncertainty std. dev.	5.0°
initial speed uncertainty std. dev.	0.1 cm/T
gate parameter γ	9
track initiation parameters	$M = 2, N = 3$

The initial position of the sonar is marked by an arrow. The end of the path is marked by a square. Once one pass through the path was completed, another one was initiated. This process was repeated until the termination of the experiment. Figure 2 (bottom right) displays the actual and estimated path of the AUV along with the estimated feature locations and their 3σ bonds (99% highest confidence regions) for the experiment. Figure 3 shows the position, heading and speed error along with the 3σ error bounds for the experiment. Figure 4 shows the distance error of the first feature initiated for the experiment. As expected, the error bound is a monotonically decreasing function of time.

When the same trajectory can be repeated multiple times, the error bounds converge to steady-state levels. The algorithm is quite successful when applied to the testing tank data set, because the amount of data association ambiguity is low, and the features are spatially distinct.

3 Forward look sonar data acquisition

A series of HRA forward look imaging sonar data sets were collected in September 1997 as part of a collaborative project between the Naval Undersea Warfare Center (NUWC) in Newport, Rhode Island, and Groupe d’Etudes Sous-Marines de l’Atlantique (GESMA) in Brest, France [4]. An illustration of the part of Narragansett Bay where the data was collected is shown in Figure 5. An over-the-side test rig, shown in Figure 6, was built to carry several typical AUV subsystems for deployment over the side of the YFRT-287, a converted U.S. Navy yard freighter. When fully deployed, the AUV subsystems

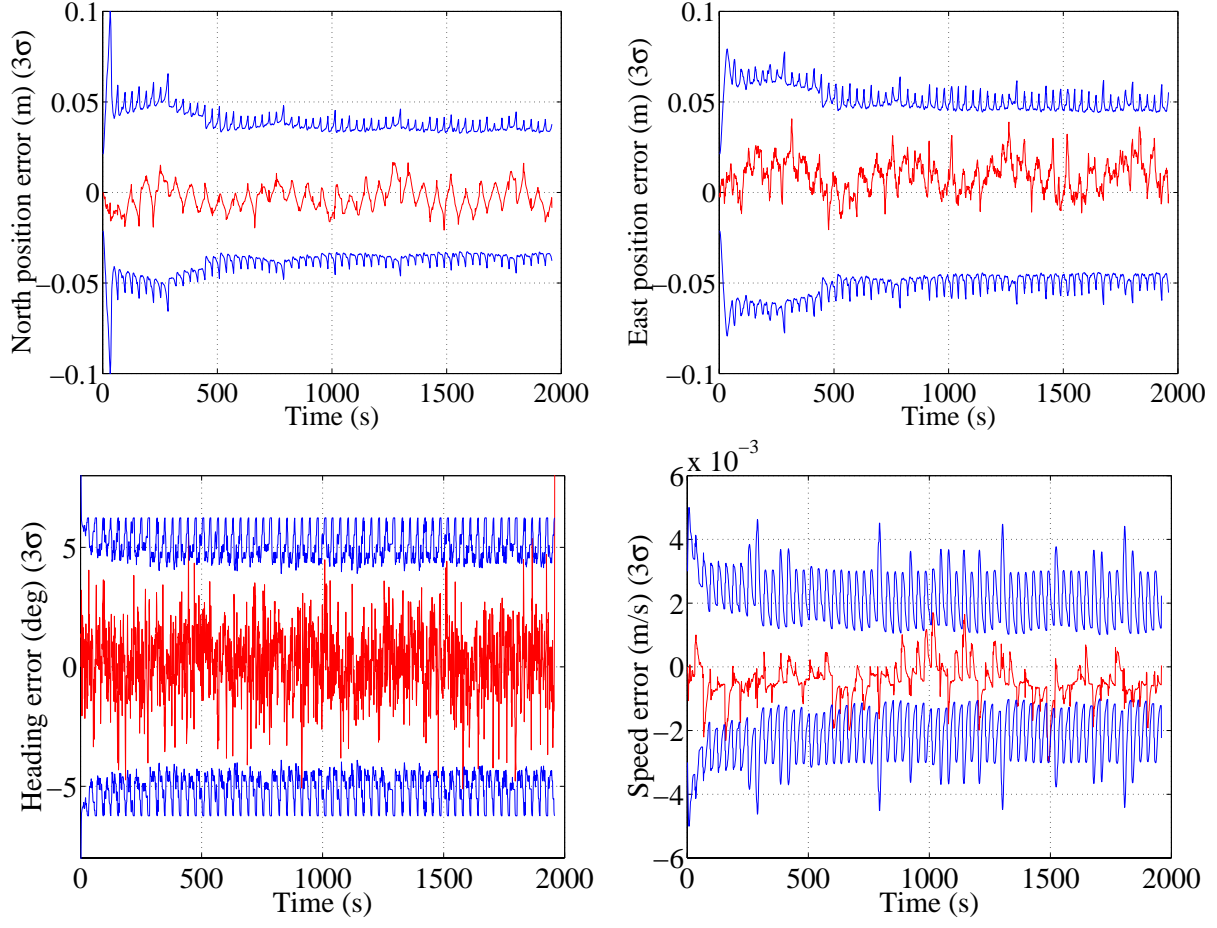


Figure 3: Errors and 3σ bounds (99% highest confidence bounds) for the position, heading and velocity estimates for the testing tank experiment.




Figure 4: (Left) Estimated feature errors and 3σ bounds. (Right) Evolution of feature estimation error and the 3σ bound versus time for the first feature initialized.

were approximately 5.5 meters below the water's surface. The water depth in the operating area varied between approximately 6 and 18 meters.

The transmit transducer used for data collection was a high frequency projector manufactured by the International Transducer Corporation (ITC), operating at 87 kHz. The receive sonar array was the high resolution array forward look sonar developed at NUWC [22]. The HRA is a forward looking planar array designed for operation in a 21-inch diameter vehicle. It consists of 1264 half-wavelength elements of 1-3 composite material (PZT-5H) configured in a 20 wavelength circular aperture (design frequency of 87 kHz). A subset of the elements consisting of two rows of 32 elements each provide single ping transmit coverage of approximately 90 degrees in azimuth by 45 degrees in elevation with a source level of approximately 205 dB.

Figure 7 shows the location of the transmit and receive elements of the HRA. The transmit-only elements are shown in black and the receive-only elements are shown in gray. Due to hardware limitations, the DAP is able to record a 0.67 second window of HRA data every 20 seconds (nominally). Due to hard-drive capacity, the DAP is also limited to 50 acoustic events before the data must be transferred to archival storage media. HRA azimuth and elevation beampatterns generated with real data are shown in Figure 8.

The HRA element analog outputs are processed with the data acquisition processor (DAP). The DAP conditions each of 511 analog signals, performs A/D conversion, and generates basebanded in-phase and quadrature data for the HRA element outputs. The DAP also records 511 channels of the element level HRA data. Beamforming is performed using the element level data to produce a complete sonar image covering from ± 40 degrees in azimuth out to 300 meters in range.

In addition to the HRA, the DAP, and the ITC transmit transducer, the following subsystems were used on the over-the-side test rig: an Allied Signal Model RL-34 Inertial Navigation System (INS), an Edo Model 3050 Doppler velocity sonar (DVS), a Trimble Model NT200D differential global positioning system (DGPS) receiver, and a Klein 2000 side scan sonar. (The Klein 2000 side scan sonar images were not processed in this paper, but provided independent verification that many bottom objects were present in the area where the data was collected.) Data from each of the subsystems were recorded during the exercise for further processing in the laboratory. Time-synchronization of the subsystems was accomplished via clock signals from the DGPS. Latitude and longitude data are converted to a locally referenced Cartesian coordinate system.

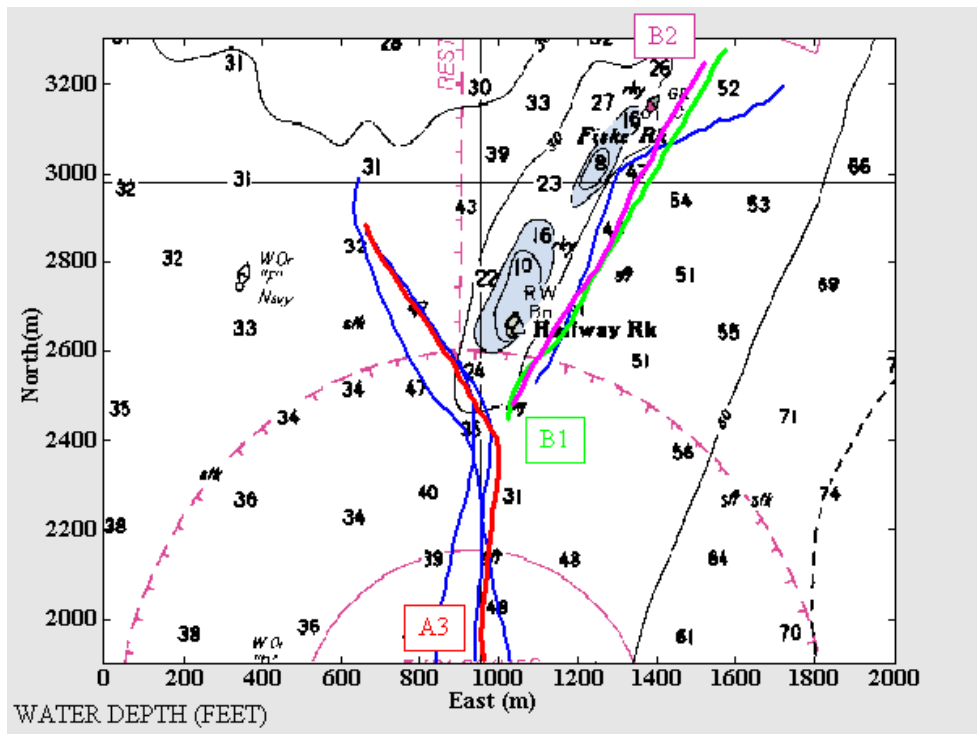



Figure 5: Narragansett Bay area around NUWC (top), and the run area (bottom). A series of data collection legs approximately 1 kilometer in length were planned and conducted in the exercise. With a 20 second pulse repetition rate and typical YFRT-287 speeds of approximately 2 knots, HRA data were collected about every 20 meters. The data sets processed in this paper were acquired on legs A1 and A3. Leg A3 is shown in the lower left of the bottom figure. Leg A1 followed a similar trajectory offset slightly to the right.




Figure 6: Over-the-side test rig. The system consisted of the NUWC HRA forward-look sonar system, a Klein 500 kHz side-scan sonar system, an Allied INS system, and an EDO Inc. Doppler velocity sonar. This system effectively constitutes an in-water “UUV simulator” which allows real-data testing of CML algorithms.




Figure 7: HRA Transmit and receive elements.

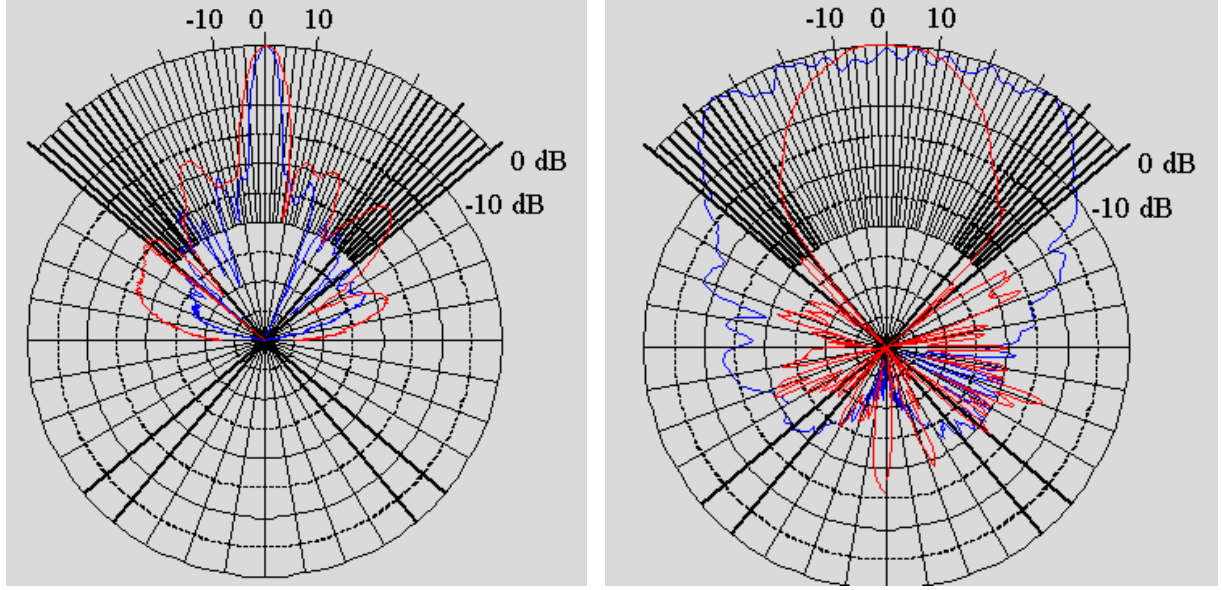


Figure 8: HRA beampatterns (Left: elevation, Right: azimuth).




Figure 9: Summary of HRA signal processing steps.

4 Feature extraction

The intent of the sonar feature extraction processing is to detect, localize and extract features from compact regions of strong bottom backscatter. Each localized region is referred to as an object. The processing sequence is summarized in Figure 9 and consists of the following steps [4]:

1. **Beamformer stage 1:** The first step in the beamformer consists of applying amplitude shading to the receive-only elements for mainlobe and sidelobe control. A set of amplitude weights is chosen to provide beams with mainlobes that are wider in the vertical than in the horizontal. This is because, in the proposed approach, beams are steered to only one elevation angle and the wider vertical beams provide better vertical coverage.

The two-dimensional matrix, \mathbf{a} , of amplitude weights are found via the outer product of two vectors of length 40, $\mathbf{a} = \mathbf{a}_v \mathbf{a}_h^T$, where \mathbf{a}_v is vector consisting of a 40 point Kaiser window ($\beta = 16$) and \mathbf{a}_h

is a vector consisting of a 40 point Hamming window. The weights are normalized to provide unity gain and the result is referred to as hybrid weights. The Kaiser window was chosen for the vertical shading because it is equipped with a design parameter β that allows for a variable beam width — an attractive feature for operation in variable water depths. The boresight beam patterns resulting from the application of the two-dimensional weights to real HRA data collected in NUWC's acoustic test facility is shown in Figure 8. The 3 dB down vertical beamwidth is approximately 7 degrees. The 3 dB down horizontal beamwidth is approximately 4 degrees.

After amplitude weighting, complex phase shading is applied to steer the elements to one elevation angle that is a function of the water depth and the altitude of the vehicle. The elevation steer angle used for all results in the paper was 5 degrees below horizontal. The elements in each column are then summed to form a set of 40 stave outputs. The sample rate, f_s , of each stave output is 10,875 complex samples per second. Typically 0.66 second ≈ 7177 samples (N_{samp}) of data are available for processing.

2. **Matched filtering:** The 40 complex stave outputs are matched filtered using a replica of the transmit signal. The current approach uses a linear frequency modulation signal with a time duration of approximately 10 ms and a bandwidth of approximately 8000 Hz.
3. **Beamformer stage 2:** The matched filter outputs are phaseshift beamformed to 128 beams in azimuth via a discrete Fourier transform that is implemented with a fast Fourier transform. Only the magnitude squared values from the outputs of the second stage of beamforming are required for subsequent processing. The data are also edited to retain only those receive beams within 3 dB down limits of the transmit beam ($\approx \pm 40$ degrees). The output from the second stage of beamforming for each listen interval consists of an $(N_{samp} \times N_{ang} = 81)$ matrix of real numbers referred to as the range-azimuth map.
4. **Normalization:** The normalization step seeks to remove propagation and environmental effects from the range-azimuth map. The fundamental approach in this effort uses statistics of reference cells that surround, and include, each test cell in the range-azimuth map. At the range increment of interest, the median value across all 81 angle cells is computed. Each and every angle cell at the range increment in question is then normalized by the computed sample median. This step is repeated at each range increment. The output from the normalization stage for each listen interval consists of an $(N_{samp} \times N_{ang})$ matrix of real numbers referred to as the normalized range-azimuth

map. Each value in the matrix represents a measure of the signal-to-noise ratio (SNR) of its corresponding cell in the range-azimuth plane.

5. **Binary image formation:** A threshold is applied to the normalized range-azimuth map to form a new map (a value corresponding to 16 dB is currently used). All cells with values equal to or greater than the threshold are given a value of 1. The rest are given a value of 0. This map is referred to as the binary-image map and has all the attributes of a binary image so that image processing techniques can be used to extract information from it.
6. **Morphological closure:** A series of morphological operations are performed on the binary image map. The intent is to group together threshold crossings from the normalized range-azimuth map to form “objects” that capture areas of localized, strong scattering from the bottom, and remain consistent in their structure and location from ping to ping. This step is also motivated by the desire to reduce the dimensionality of the data in the binary-image map.

The primary morphological operation used in the planar processing algorithm is closure. Two closure operations are performed in the planar processing algorithm with two different structuring elements (SE). The first SE measures 14 range cells (≈ 1 meter) by 1 angle cell, and the second measures 72 range cells (≈ 5 meters) by one angle cell.

7. **Ordered perimeters:**

The perimeter of each object is computed via a series of morphological operations. First, a dilation is performed on the object to remove any “spurs” (the SE is a 3 by 3 matrix of ones). Next “holes” in objects are filled in. The perimeter is then calculated by subtracting an eroded version of this object from itself. The erosion in this step uses a 3 by 3 matrix of ones as the SE. The perimeter of each object is then ordered so that when perimeter pixels are connected, the line segments do not cross. The perimeters of each object are then converted to two-dimensional Cartesian coordinates using the center of the array as the origin. As noted previously, vehicle heading is incorporated to provide an absolute reference to north. The small roll and pitch deviations of the vehicle are also used in the conversion from polar to rectangular coordinates. The result is that each object is represented as a possibly-nonconvex polygon.

The object perimeters for six sonar pings from data set A3 are shown in Figure 10. The range and bearing to the centroid of each object is computed and represents the measurement that is applied to the stochastic mapping algorithm.

Our method for CML models the environment in terms of point-like features. This assumption clearly does not hold for some of the features in the environment, such as the ridge encountered midway through legs A1 and A3. Large objects corresponding to the ridge are visible in the data from pings 23 and 24 in Figure 10). Because the object centroids are not as stable for large objects as for small objects, very large objects (those with an area greater than 100 meters) were screened out as part of the data association process. This corresponds to approximately 10 % of the measurements for each data set.

5 Forward look sonar stochastic mapping results

This section describes the application of the method presented in Section 2 to the HRA data. Heading and speed measurements from the INS and Doppler sensors mounted on the over-the-side test rig were used as proprioceptive sensor measurements for the stochastic mapping algorithm. Roll, pitch, and depth measurements from the test rig were incorporated in the sonar feature extraction processing, providing three-dimensional feature location estimates. Because the depth of the transducer is fixed, we restrict our attention to testing the feasibility of CML using two-dimensional models as employed in the testing tank experiment shown in Figure 2.

The HRA sonar data sets are obtained approximately every 20 seconds, due to the storage capacity of the data acquisition system. For an implementation on-board a real AUV, it is anticipated that sampling periods of one second or less can be achieved with commercial, off-the-shelf hardware. Increased sampling rates would likely produce a slower rate of growth of errors.

In an implementation of CML operating on-board a real AUV, knowledge of the vehicle’s control input would likely be available; however, for the HRA data collection exercise, there was no way to record the control inputs of the ship and a dynamic model for the ship was unavailable. For the results presented in this paper, the control inputs were assumed to be zero, representing straight-line motion at a constant speed.

The parameters for the application of the stochastic mapping algorithm to the HRA sonar data sets are shown in Table 2. We consider processing of two of the legs, A1 and A3. These were acquired on different days following very similar trajectories (as shown in Figure 5).

Figures 11 to 14 show the results for leg A1 with track initiation parameters $M = 2$ and $N = 3$. (A minimum of two matching features out of three time steps are required to initialize a new feature.) Figure 11 shows all the returns from the data set. For the data collection exercise, three strong reflecting targets (steel spheres) were placed at known locations. These are marked by ‘×’, located near the final

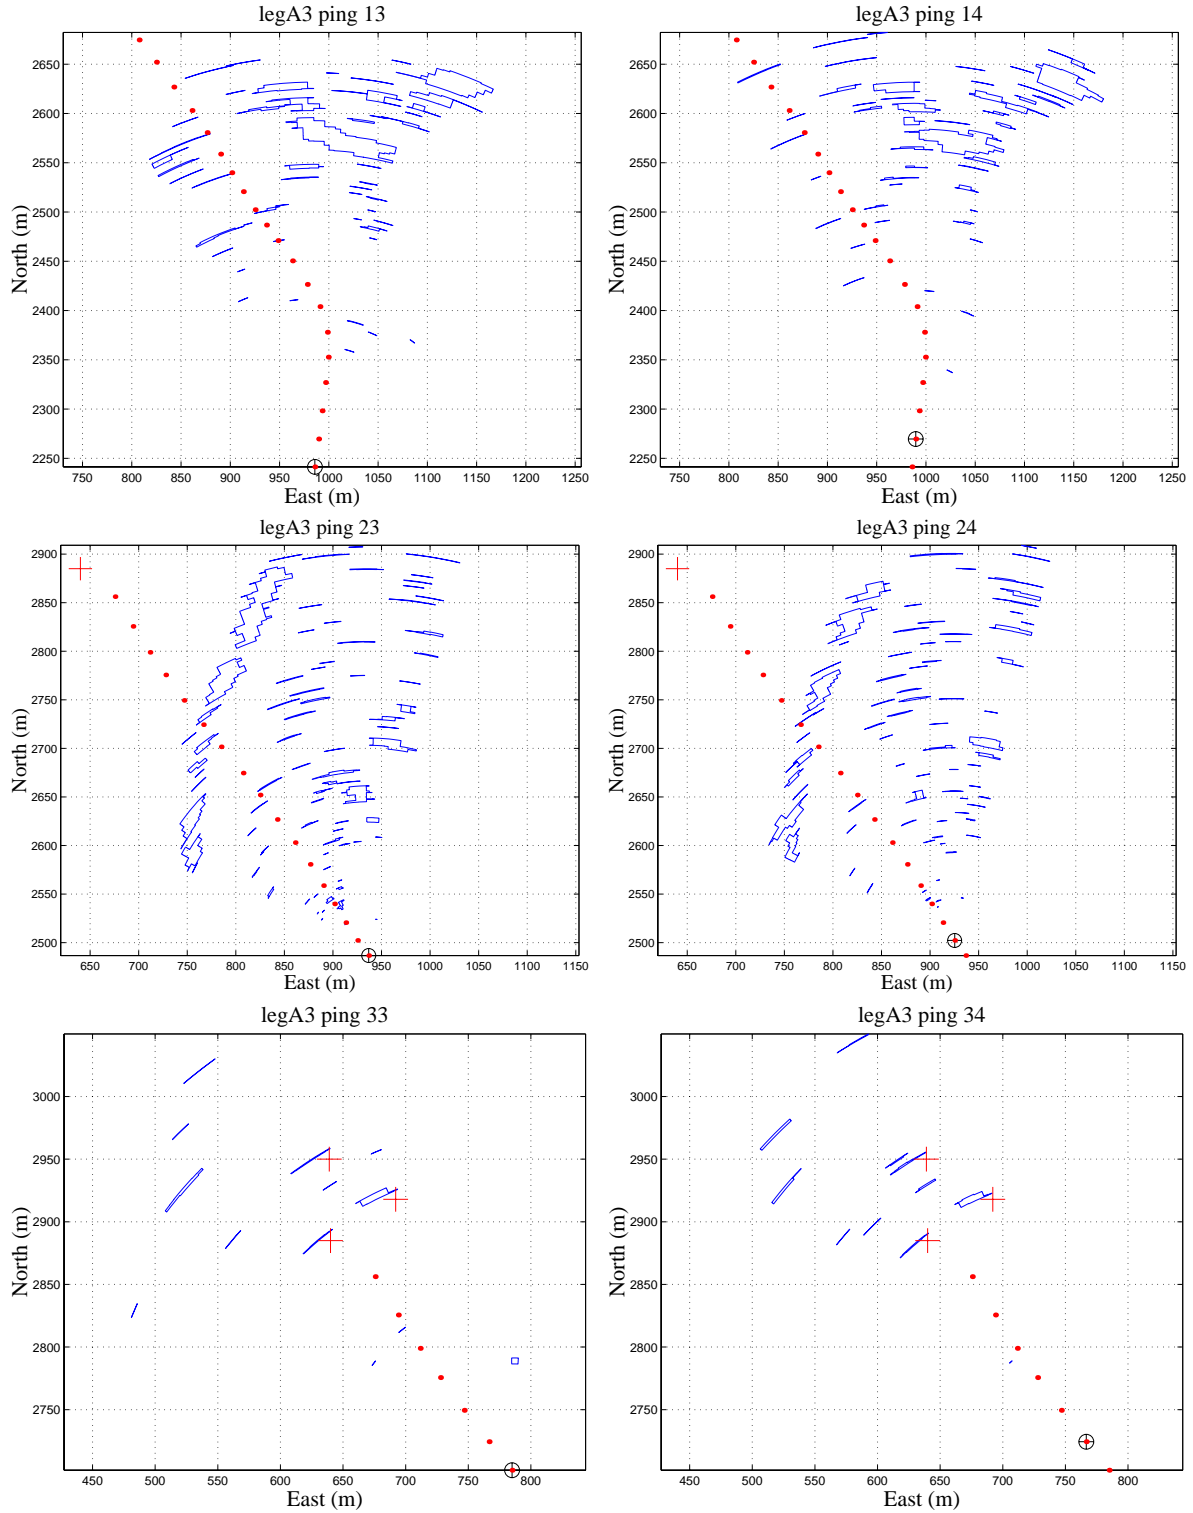


Figure 10: Object polygons for several pings from Leg A3. Small dots indicate the vehicle trajectory as determined by DGPS. The vehicle position for each ping is shown with a '+' surrounded by a circle at the bottom edge of each plot. (Top row) pings 13 and 14. (Middle row) pings 23 and 24. (Note the very large polygons corresponding to the ridge on the left side for each of these pings.) (Bottom row) pings 33 and 34. The locations of three known spherical targets are indicated with '+' marks.

Table 2: HRA data post-processing parameters.

sampling period, T	20 sec.
(x, y) process noise std. dev.	2 cm
speed process noise std. dev.	4 cm/sec
heading process noise std. dev.	3°
speed measurement std. dev.	1.0 m/sec
heading measurement std. dev.	3°
sonar range measurement std. dev.	1 m
sonar bearing measurement std. dev.	0.5°
gate parameter γ	10

position of the vehicle. In addition, there were many environmental features on the seabed, such as rocks and lobster pots. In the middle of the vehicle’s path, there was a very prominent ridge on the bottom, which produced a large amount of returns.

During the data acquisition exercise, the INS and Doppler sensor measurements were processed by a navigation Kalman filter that was designed for a Navy unmanned underwater vehicle [19]. The position error growth from this filter was approximately 2% of distance traveled.

Figure 12 (Left) displays the vehicle position estimates from DGPS, CML and INS, along with the 3σ error ellipses for the modeled environmental features. Figure 12 (Right) plots the error in the CML position estimate as a function of time, using DGPS as ground truth. The errors from CML grow more slowly than the errors obtained with INS and Doppler measurements only. Figure 13 describes the data association decisions that were made for each measurement at each time step. A total of 1099 measurements were processed. 77 features were initialized. Of these 77 features, 28 features received updates from a total of 82 measurements.

Figure 14 shows the errors and 3σ error bounds for the position, heading and velocity estimates. The error bounds increase towards the end of the run as the number of measurements decreases. The error bounds produced by a CML algorithm will grow without bound if the vehicle follows a straight line trajectory without looping back to permit re-observation of features detected earlier in the mission. With a very high sampling rate, it may be possible to make this error growth rate very small. As illustrated in the tank experiment shown in Figure 2, when the vehicle can revisit all parts of the environment multiple times, bounded errors can be achieved.

Figure 15 shows some of the measurements used to initialize and update the tracks for several of the features that obtained the highest number of measurements during this run.

Figure 16 shows the results for processing of the data from leg A3 with three different sets of track

initiation parameters. The value of M , the minimum number of returns needed to initialize a new feature, is varied from two to four. N , the set of time steps over which features are considered for track initiation, is set to $M + 1$. For $M = 2$, 122 features were initialized and 231 measurements were used for feature updates. For $M = 3$, 60 features were initialized and 132 measurements were used for feature updates. For $M = 4$, 39 features were initialized and 83 measurements were used for feature updates. Depending on the number of features that are initialized, the vehicle trajectory varies somewhat, however, the overall algorithm performance is fairly consistent.

In comparing the CML and INS position estimation errors shown above in Figure 12 and 16, it is important to note that the INS filter was not specifically developed for the over-the-side test rig system used in these tests. In actual operation on-board the Navy unmanned underwater vehicle for which the INS filter was designed, error growth rates of less than .2% were achieved. The 2% error growth rate shown in Figure 12 is more representative of the navigation performance that would be obtained with low-cost proprioceptive sensors, such as a fluxgate compass. The comparison is useful, none-the-less, to show the potential for CML to achieve a bounded error in situations where INS alone suffers from unbounded growth.

6 Conclusion

This paper has considered the problem of concurrent mapping and localization with forward look sonar data using a feature-based stochastic mapping algorithm. Processing of data from a testing tank experiment illustrates the potential of CML to achieve a bounded error in environments with distinctive point features and a low degree of clutter. Results from post-processing of an oceanic forward look sonar data set demonstrate the potential to successfully perform CML on-board an AUV. Salient features can be detected and tracked to provide effective navigation in a shallow-water operating environment.

The results indicate that more work is necessary in the areas of feature modeling and data association. For extended objects, such as the ridge on the ocean bottom that is encountered midway through the data sets, the representation of the environment as discrete points breaks down. Many point features are initiated, and the data association ambiguity is very high. Successful performance is achieved because the nearest neighbor gating strategy is very conservative. If the origin of a measurement is ambiguous, it is better to discard it, rather than to make a wrong association. The results in this paper used only the centroid of each detected object for data association. In related research, we have obtained promising initial results using other feature attributes, such as area, perimeter, and radial signature [4].

Work in progress is creating a new software implementation of CML suitable for real-time operation on-board a US Navy unmanned underwater vehicle. The creation of a real-time implementation of CML is challenging because the complexity of stochastic mapping is $\mathcal{O}(n^2)$ [21]. In the data sets processed above, up to 100 features were initialized from less than fifty sonar pings, for approximately 1 kilometer of vehicle travel. For vehicle missions of any substantial duration, it will become impossible to maintain all the correlation terms for all features in a single error covariance matrix. To address this issue, we have developed decoupled stochastic mapping, a computationally efficient method for large-scale CML [18].

Acknowledgments

This research has been funded in part by the Henry L. and Grace Doherty Assistant Professorship in Ocean Utilization, NSF Career Award BES-9733040, the MIT Sea Grant College Program under grant NA86RG0074 (project RCM-3), and the US Navy International Programs Office. The authors acknowledge the contributions of a number of people for their data collection efforts. Included are R. Ayers, C. Brown, S. Camara, D. French, V. Greco, F. Hoffman, M. Medeiros, C. Shaw, R. Shell, and J. Vaillancourt.

References

- [1] N. Ayache and O. Faugeras. Maintaining representations of the environment of a mobile robot. *IEEE Trans. Robotics and Automation*, 5(6):804–819, 1989.
- [2] Y. Bar-Shalom and T. E. Fortmann. *Tracking and Data Association*. Academic Press, 1988.
- [3] R. A. Brooks. Aspects of mobile robot visual map making. In *Second Int. Symp. Robotics Research*, Tokyo, Japan, 1984. MIT Press.
- [4] R. N. Carpenter. Concurrent mapping and localization using forward look sonar. In *IEEE AUV*, Cambridge, MA, August 1998.
- [5] J. A. Castellanos, J. M. M. Montiel, J. Neira, and J. D. Tardos. The SPmap: A probabilistic framework for simultaneous localization and map building. *IEEE Trans. Robotics and Automation*, 15(5):948–952, 1999.
- [6] R. Chatila and J. Laumond. Position referencing and consistent world modeling for mobile robots. In *IEEE International Conference on Robotics and Automation*. IEEE, 1985.
- [7] I. J. Cox and J. J. Leonard. Modeling a dynamic environment using a Bayesian multiple hypothesis approach. *Artificial Intelligence*, 66(2):311–344, April 1994.
- [8] T. Curtin, J. G. Bellingham, J. Catipovic, and D. Webb. Autonomous ocean sampling networks. *Oceanography*, 6(3):86–94, 1993.
- [9] M. W. M. G. Dissanayake, P. Newman, H. F. Durrant-Whyte, S. Clark, and M. Csorba. A solution to the simultaneous localization and map building (SLAM) problem. In *Sixth International Symposium on Experimental Robotics*, March 1999.

- [10] H. F. Durrant-Whyte, M. W. M. G. Dissanayake, and P. Gibbons. Towards simultaneous localization and map building (SLAM) in large unstructured environments,. In D. Koditschek and J. Hollerbach, editors, *Robotics Research: The Ninth International Symposium*, pages 162–168, Snowbird, Utah, October 1999. Springer Verlag.
- [11] H. J. S. Feder. *Simultaneous Stochastic Mapping and Localization*. PhD thesis, Massachusetts Institute of Technology, 1999.
- [12] H. J. S. Feder, J. J. Leonard, and C. M. Smith. Adaptive mobile robot navigation and mapping. *Int. J. Robotics Research*, 18(7):650–668, July 1999.
- [13] H. J. S. Feder, C. M. Smith, and J. J. Leonard. Incorporating environmental measurements in navigation. In *IEEE AUV*, Cambridge, MA, August 1998.
- [14] E. Geyer, P. Creamer, J. D’Appolito, and R. Gains. Characteristics and capabilities of navigation systems for unmanned untethered submersibles. In *Proc. Int. Symp. on Unmanned Untethered Submersible Technology*, pages 320–347, 1987.
- [15] J.-S. Gutmann and K. Konolige. Incremental mapping of large cyclic environments. In *Proc. IEEE Int. Conf. Robotics and Automation*, 2000.
- [16] J. J. Leonard, A. A. Bennett, C. M. Smith, and H. J. S. Feder. Autonomous underwater vehicle navigation. Technical Report Marine Robotics Laboratory Technical Memorandum 98-1, MIT, 1998.
- [17] J. J. Leonard and H. F. Durrant-Whyte. *Directed Sonar Sensing for Mobile Robot Navigation*. Boston: Kluwer Academic Publishers, 1992.
- [18] J. J. Leonard and H. J. S. Feder. A computationally efficient method for large-scale concurrent mapping and localization. In D. Koditschek and J. Hollerbach, editors, *Robotics Research: The Ninth International Symposium*, pages 169–176, Snowbird, Utah, October 1999. Springer Verlag.
- [19] E. Levine, D. Connors, R. Shell, T. Gagliardi, and R. Hanson. Oceanographic mapping with the Navy’s large diameter UUV. *Sea Technology*, pages 49–57, 1995.
- [20] M. Medeiros and R. Carpenter. High resolution array signal processing for AUVs. In *AUV 96*, pages 10–15, 1996.
- [21] P. Moutarlier and R. Chatila. Stochastic multisensory data fusion for mobile robot location and environment modeling. In *5th Int. Symposium on Robotics Research*, Tokyo, 1989.
- [22] F. Nussbaum, G. Stevens, and J. Kelly. Sensors for a forward-looking high resolution AUV sonar. In *AUV 96*, pages 141–145, 1996.
- [23] D. B. Reid. An algorithm for tracking multiple targets. *IEEE Trans. on Automatic Control*, AC-24(6), Dec. 1979.
- [24] M. Skolnick. *Introduction to Radar Systems*. New York: McGraw-Hill, 1981.
- [25] C. M. Smith. *Integrating Mapping and Navigation*. PhD thesis, Massachusetts Institute of Technology, 1998.
- [26] C. M. Smith, H. J. S. Feder, and J. J. Leonard. Multiple target tracking with navigation uncertainty. In *IEEE Int. Conference on Decision and Control (CDC)*, Tampa, FL, December 1998.
- [27] R. Smith, M. Self, and P. Cheeseman. A stochastic map for uncertain spatial relationships. In *4th International Symposium on Robotics Research*. MIT Press, 1987.
- [28] R. Smith, M. Self, and P. Cheeseman. Estimating uncertain spatial relationships in robotics. In I. Cox and G. Wilfong, editors, *Autonomous Robot Vehicles*. Springer-Verlag, 1990.
- [29] W. K. Stewart. *Multisensor Modeling Underwater with Uncertain Information*. PhD thesis, Massachusetts Institute of Technology, 1988.
- [30] S. Thrun, D. Fox, and W. Burgard. A probabilistic approach to concurrent mapping and localization for mobile robots. *Machine Learning*, 31:29–53, 1998.




Figure 11: Stochastic mapping post-processing results for leg A1 with track initiation parameters $M = 2, N = 3$. Returns from the HRA sonar are shown by small dots. The DGPS position fixes at each measurement point are shown by triangles. The CML result is drawn as a solid line. The output from the INS/Doppler navigation filter running on the over-the-side test rig during the experiment is shown by small circles. The three known features are represented by ‘ \times ’ signs, and the estimated feature locations are shown by ‘+’ signs. (Top) The result for the entire mission. (Bottom) Magnified view of the state estimates at the end of the mission.

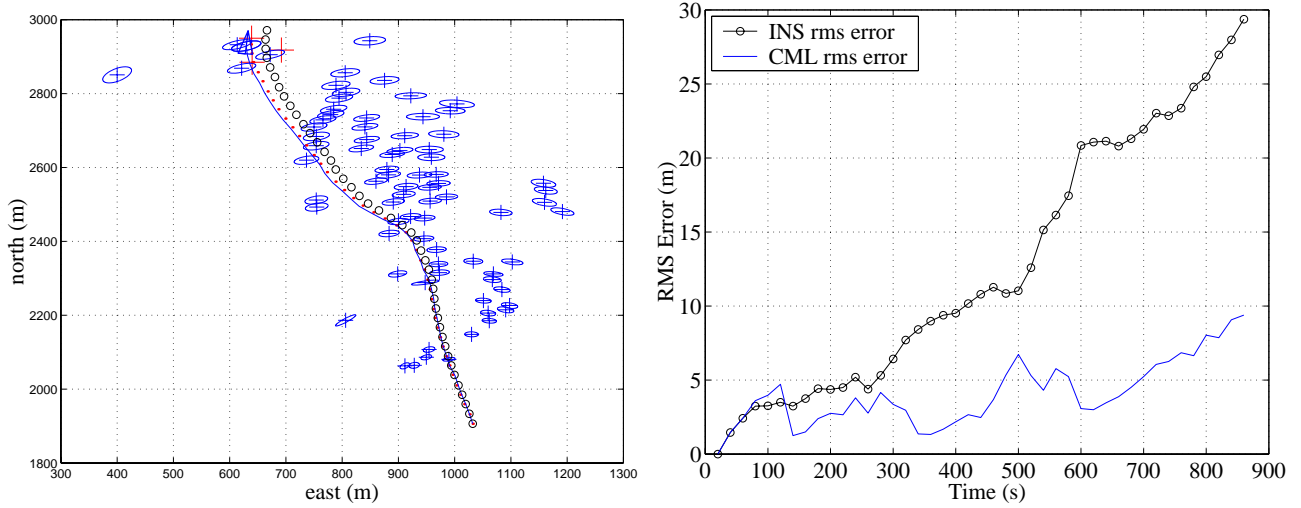


Figure 12: (Left) The estimated feature locations, marked by '+', along with the 3σ ellipses for each of these estimates for post-processing of the data from leg A1. (Right) A comparison of root mean square error when relying on the inertial navigation system only versus the CML result produced by stochastic mapping as a function of time. (DGPS data is used for ground-truth.)




Figure 13: The number of returns from the HRA sonar that were used for track updating and track initiation at each time step. Note that only a very small number of returns were used for track updating and thus for improving mapping and localization in the environment. With the data association strategy employed in stochastic mapping, many returns are discarded because their origin is considered too ambiguous. However, improved navigation is realized with only a small number of correctly associated measurements.




Figure 14: Errors and 3σ bounds (99% highest confidence bounds) for the position, heading and velocity estimates. (DGPS data is used for ground-truth.)

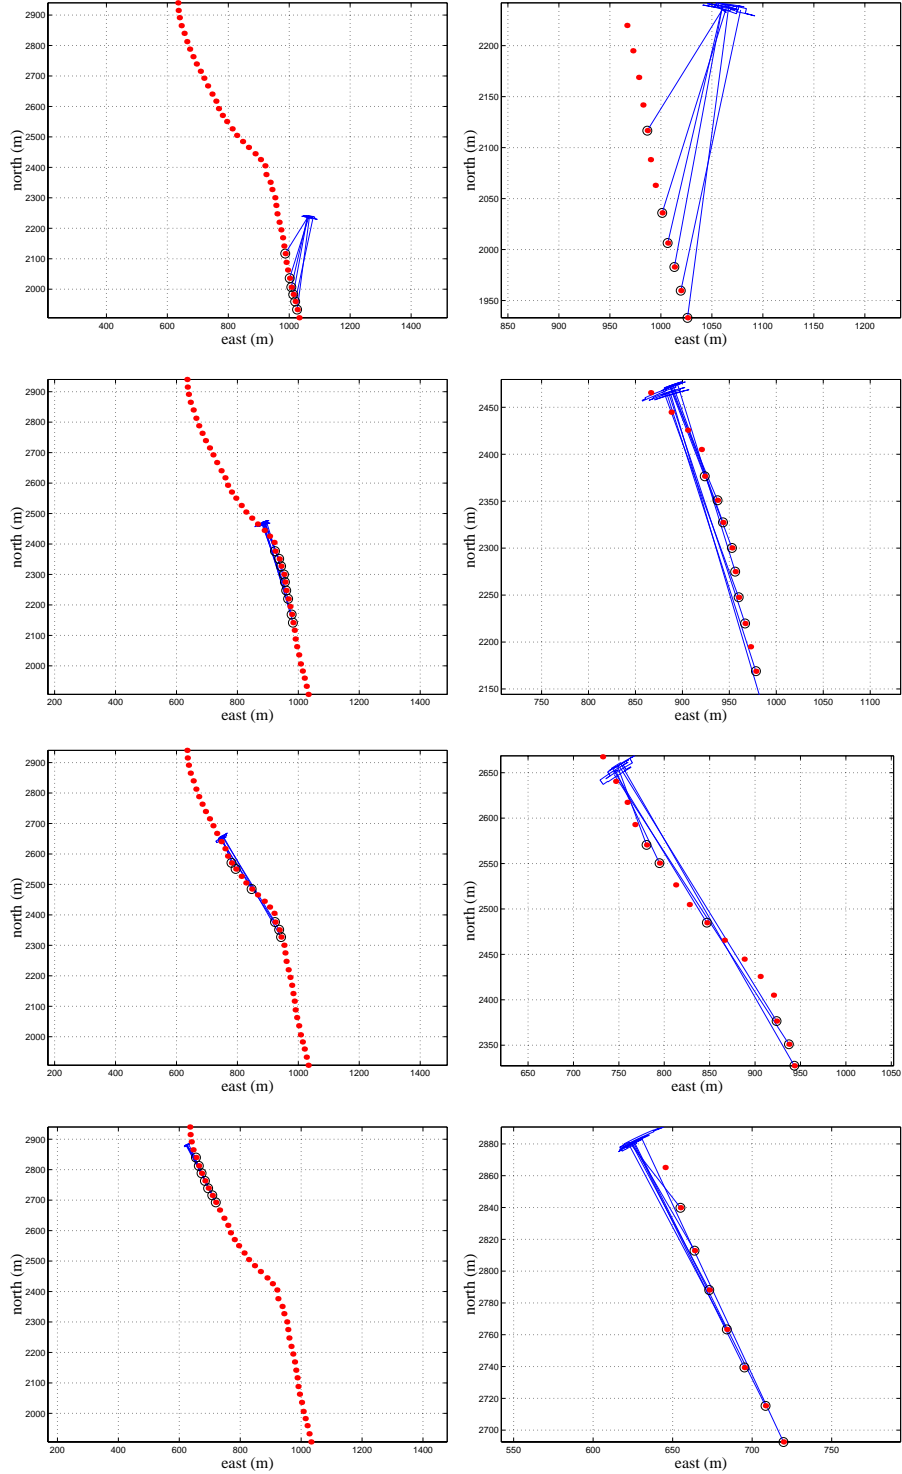


Figure 15: Measurements associated with several different features during processing of leg A1. The complete vehicle trajectory estimated by CML is shown as a sequence of dots. The right column shows a magnified view of the measurements.

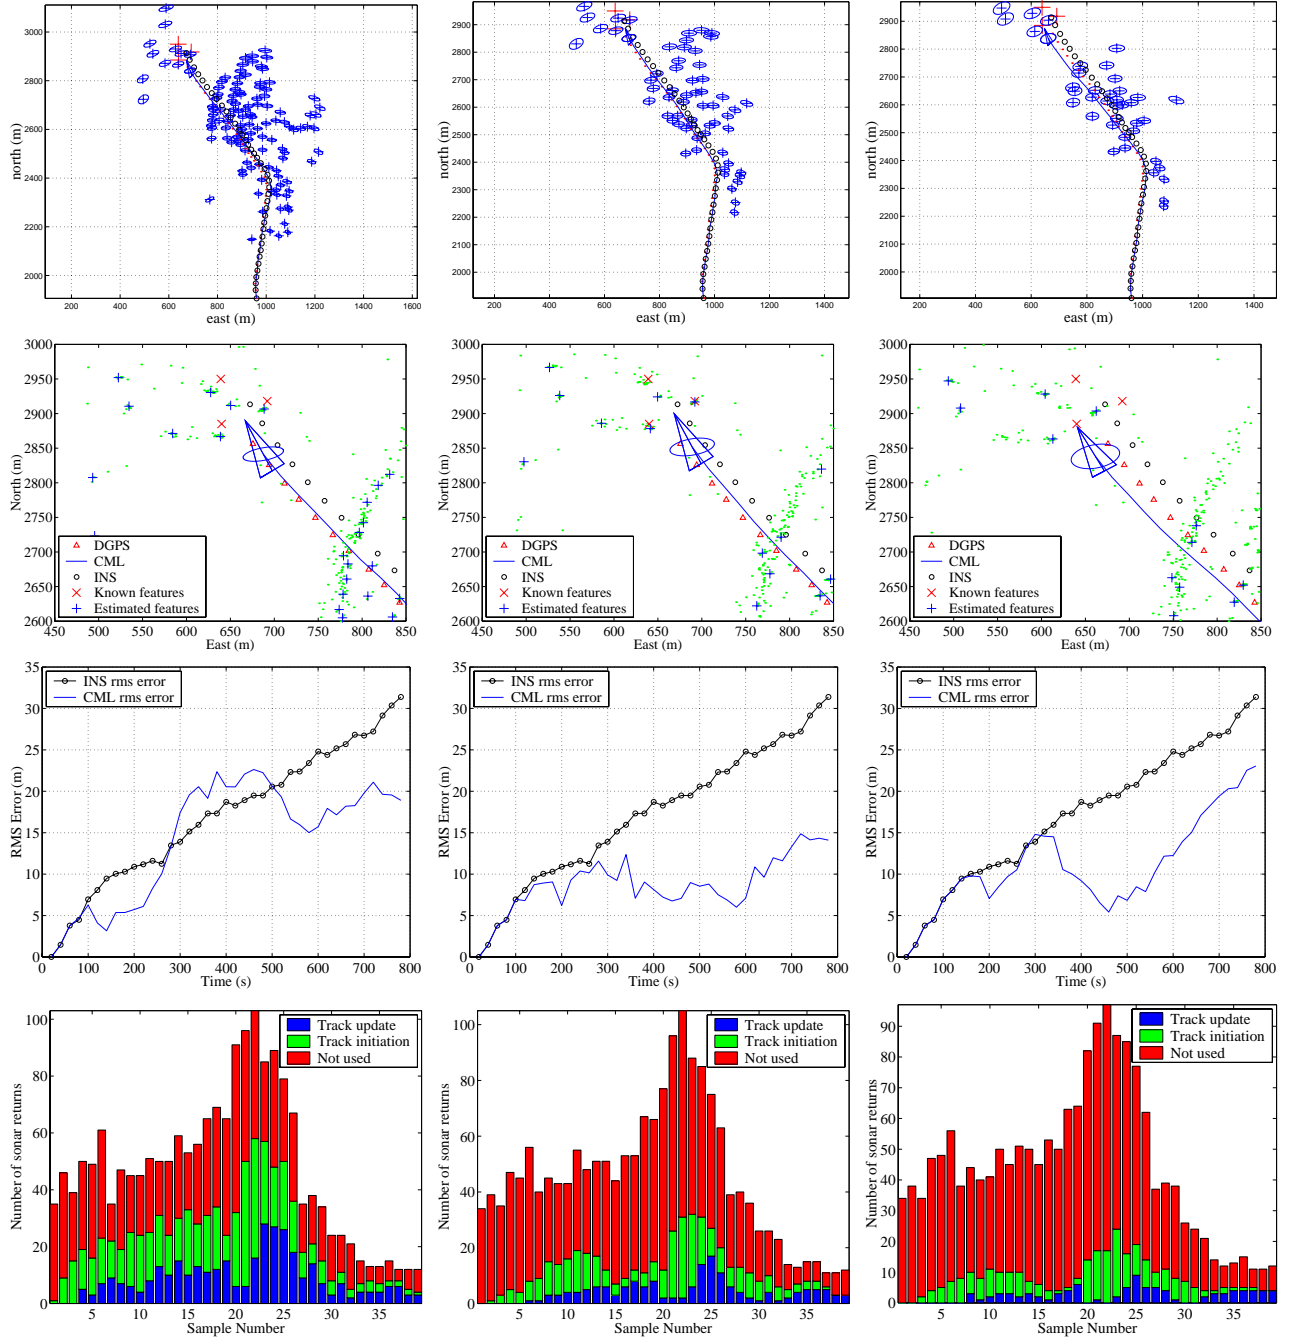


Figure 16: Results for leg A3 with three different values for track initiation parameters. (Left column) $M = 2, N = 3$, (Middle column) $M = 3, N = 4$, (Right column) $M = 4, N = 5$. The plots on the top row show the estimated feature locations. The plots on the second row show a magnified view of the CML, DGPS, and INS state estimates at the end of the mission. The plots on the third row compare the root mean square error when using INS/Doppler only with the CML result. Finally, the plots on the bottom row show the numbers of returns used for track initiation and track updating.